



A Novel Evaluation Perspective on GNNs-based Recommender Systems through the Topology of the User-Item Graph

Daniele Malitesta*[†]
 Université Paris-Saclay
 CentraleSupélec, Inria
 Gif-sur-Yvette, France
 daniele.malitesta@centralesupelec.fr

Claudio Pomo*
 Politecnico di Bari
 Bari, Italy
 claudio.pomo@poliba.it

Vito Walter Anelli
 Politecnico di Bari
 Bari, Italy
 vitowalter.aneli@poliba.it

Alberto Carlo Maria Mancino
 Politecnico di Bari
 Bari, Italy
 alberto.mancino@poliba.it

Tommaso Di Noia
 Politecnico di Bari
 Bari, Italy
 tommaso.dinoia@poliba.it

Eugenio Di Sciascio
 Politecnico di Bari
 Bari, Italy
 eugenio.disciascio@poliba.it

ABSTRACT

Recently, graph neural networks (GNNs)-based recommender systems have encountered great success in recommendation. As the number of GNNs approaches rises, some works have started questioning the theoretical and empirical reasons behind their superior performance. Nevertheless, this investigation still disregards that GNNs treat the recommendation data as a topological graph structure. Building on this assumption, in this work, we provide a novel evaluation perspective on GNNs-based recommendation, which investigates the impact of the graph topology on the recommendation performance. To this end, we select some (topological) properties of the recommendation data and three GNNs-based recommender systems (i.e., LightGCN, DGCF, and SVD-GCN). Then, starting from three popular recommendation datasets (i.e., Yelp2018, Gowalla, and Amazon-Book) we sample them to obtain 1,800 size-reduced datasets that still resemble the original ones but can encompass a wider range of topological structures. We use this procedure to build a large pool of samples for which data characteristics and recommendation performance of the selected GNNs models are measured. Through an explanatory framework, we find strong correspondences between graph topology and GNNs performance, offering a novel evaluation perspective on these models.

CCS CONCEPTS

• Information systems → Personalization.

KEYWORDS

Graph Neural Networks, Topology, Recommender Systems

*Corresponding authors.

[†]Work done while at Politecnico di Bari as a PhD student.



This work is licensed under a Creative Commons Attribution International 4.0 License.

RecSys '24, October 14–18, 2024, Bari, Italy
 © 2024 Copyright held by the owner/author(s).
 ACM ISBN 979-8-4007-0505-2/24/10
<https://doi.org/10.1145/3640457.3688070>

ACM Reference Format:

Daniele Malitesta, Claudio Pomo, Vito Walter Anelli, Alberto Carlo Maria Mancino, Tommaso Di Noia, and Eugenio Di Sciascio. 2024. A Novel Evaluation Perspective on GNNs-based Recommender Systems through the Topology of the User-Item Graph. In *18th ACM Conference on Recommender Systems (RecSys '24)*, October 14–18, 2024, Bari, Italy. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3640457.3688070>

1 INTRODUCTION AND MOTIVATIONS

Graph neural networks (GNNs)-based recommender systems [39, 66] are regarded as well-established approaches within the plethora of existing recommendation algorithms, together with traditional model families encompassing, among others, neighborhood-based solutions [43, 50], latent-factor approaches [8, 19, 46, 47], as well as more recent frameworks based, for instance, on deep neural networks [17, 27], variational autoencoders [10, 52], and generative models [26, 57].

In early years, GNNs-based recommender systems proposed to design and apply GNN message-passing layers from the graph representation learning domain (i.e., graph convolutional layers such as GraphSAGE [23] and GCN [30]) to tailor them for the recommendation task. In this respect, pioneer works such as [55, 58, 70] put the selected GNN layer on top of the traditional latent-factor recommendation module; that is, users' and items' embeddings are initially refined through the message-passing procedure that can catch existing user-item relationships at multiple hops.

Over the years, other solutions tried to improve the existing techniques in various manners, inspired by the advances in graph representation learning. Among them, some approaches proposed to simplify [64] the message-passing procedure through the removal of feature transformation and non-linearities [14, 25], while others exploited the graph edges to weight their importance [6, 40], in some cases through attention weights (as in the original GAT [56] architecture) to prune noisy interactions and disentangle them at the granularity of intents [36, 59, 60, 73].

More recently, novel trends involved semi-supervised [28] and contrastive learning [29] to tackle the data sparsity of the user-item graph through ad-hoc graph augmentations [11, 33, 34, 65, 71, 72]. In parallel, other techniques suggested that addressing foundational problems in graph representation learning (such as over-smoothing [13], over-squashing [54], or both phenomena [21])

could also improve GNNs-based recommendation by surpassing the traditional concept of message-passing [41, 44, 51, 67].

As new GNN solutions are continuously proposed in the literature, a fundamental question arises: “*Why GNNs-based recommendation works so well?*”. Indeed, the dilemma is pertinent since more and more techniques are introduced daily, and knowing which strategies may better suit the recommendation problem (and why) has become imperative. In this respect, the literature recognizes that two complementary research directions have emerged so far.

On the one side, some works investigated the *theoretical* reasons behind the performance of GNNs-based recommender systems. Among others, Shen et al. [51] observed that the graph convolution operation could be behind the success of GNNs-based recommendation, working as a smoothing filter on nodes features; by leveraging the theory of graph signal processing, the authors revisited several latent factors and GNNs-based models under a shared framework built upon the smoothing filtering operation, eventually proposing a lightweight closed-form solution that can reach the performance of other less shallow recommendation solutions. More recently, Wu et al. [67] discussed the role of over-smoothing and over-correlation in GNNs-based recommendation; by acknowledging that the two phenomena are intrinsically bound, the authors proposed an adaptive de-correlation loss function to tackle the two issues at once.

On the other side, some works explored the *experimental* reasons behind the performance of GNNs-based recommender systems. For instance, the work by Wang et al. [62] tried to generalize the GNN recommendation pipeline into separate building blocks that can be changed and removed accordingly; by exploring the numerous different architectures in the literature, the authors defined a search space that could easily be adapted for every recommendation scenario. Then, the works in [4, 5] provided an outlook on accuracy and beyond-accuracy recommendation performance of GNNs-based recommender systems, recognizing node representation and neighborhood exploration as the key strategy patterns to build any recommendation model based upon message-passing. More recently, the work by Anelli et al. [7] proposed a reproducibility analysis on popular GNNs-based recommender systems spanning several existing solutions; while confirming the effective reproducibility of such models for the experimental settings presented in the original papers, the authors outlined how the commonly observed performance may drastically change on usually untested settings (i.e. other baselines or datasets), further suggesting that dataset intrinsic properties (such as the average node degree) may be greatly influencing the final performance.

Nevertheless, we believe both outlined research directions still disregard one simple (but potentially fundamental) aspect. Since GNNs are designed to treat the recommendation data as a bipartite and undirected graph, there is a whole set of *topological structures* of the user-item graph that such models should be able to recognize and exploit to boost the recommendation performance. Thus, answering the question: “*Why GNNs-based recommendation works so well?*”, would eventually imply answering another question: “*Which topological dataset characteristics are GNNs-based recommender systems able to capture? Are those properties influencing the recommendation performance?*”.

Noticeably, our intuition finds supporting motivations in the recent literature. On the one hand, works in graph representation

learning have started questioning what is the role of the graph structure (both topology and node features) in the performance of GNNs models [1, 12, 31, 31, 35, 48, 63, 64, 68, 69, 74]; for instance, it is currently acknowledged that GNNs excel in node classification tasks when working on graphs exhibiting high levels of *homophily* [75]; similarly, recent work suggested that node degree can greatly influence the performance of GNNs in both node classification and link prediction tasks [61]. On the other hand, existing studies in recommendation outlined that dataset characteristics (such as the dataset sparsity) may impact the final performance [2, 18, 37].

1.1 Our contributions

For all the outlined aspects, in this work, we propose a novel evaluation perspective on GNNs-based recommender systems, that accounts for the *topological properties* of the user-item graph. While our paper ideally settles in between the *theoretical* and *experimental* research lines outlined above, a fundamental starting basis is provided by the works in [7, 38], where the authors eventually discuss how node degree information (propagated at multiple hops) may affect recommendation accuracy for different users’ groups in the system. We decide to extend their reproducibility and experimental settings by proposing a *novel evaluation pipeline* for GNNs-based recommender systems, where we seek to find statistical correspondences between the topological properties of the user-item graph and the recommendation performance of such models. Thus, our contributions are summarized as:

- (1) We select five *classical* dataset characteristics (presented in [2, 18]) and three additional *topological* dataset characteristics [32, 42] (node degree, clustering coefficient, and degree assortativity); for the *topological* ones, we provide a re-interpretation under the lens of recommendation.
- (2) We consider three popular and recent GNNs-based recommendation approaches (i.e., LightGCN [25], DGCF [60], and SVD-GCN [44]) and work on their formulations to explicitly highlight the presence of the *topological* properties.
- (3) We propose a novel evaluation pipeline that investigates the influence of dataset characteristics on the recommendation performance of GNNs-based recommender systems; this involves the adoption of ad-hoc graph sampling strategies to generate a large suite of 1,800 small datasets that resemble three popular recommendation datasets (i.e., Yelp2018, Gowalla, and Amazon-Book) but encompass a wider range of topological structures.
- (4) Through an explanatory framework, we show that strong correspondences exist between dataset characteristics and recommendation performance for GNNs-based models, offering a novel evaluation perspective on such approaches.
- (5) To assess the goodness of the evaluation pipeline, we analyze the possible impact of node- and edge-dropout on the generated explanations, proving that their joint adoption is beneficial to the statistical significance of the explanatory framework.

To foster reproducibility, we release the code to replicate our evaluation pipeline: <https://github.com/sisinflab/Topology-Graph-Collaborative-Filtering>.

2 TOPOLOGICAL DATA CHARACTERISTICS

In this section, we provide a formal description of the recommendation dataset in terms of its *topological* characteristics [32, 42], treating the data as a bipartite and undirected graph; noticeably, for each topological property, we also offer a re-interpretation under the lens of recommendation. As our analysis will eventually assess also the influence of *classical* recommendation dataset characteristics already presented in [2, 18], and for the sake of space, we suggest the reader refer to those works for a formal presentation.

In a recommendation system, we denote with \mathcal{U} and \mathcal{I} the sets of users and items, respectively, where $|\mathcal{U}| = U$ and $|\mathcal{I}| = I$. Then, we indicate with $\mathbf{R} \in \mathbb{R}^{U \times I}$ the interaction matrix collecting user-item interactions in the form of implicit feedback (i.e., $\mathbf{R}_{u,i} = 1$ if user $u \in \mathcal{U}$ interacted with item $i \in \mathcal{I}$, 0 otherwise).

The interaction matrix \mathbf{R} is leveraged to define the adjacency matrix $\mathbf{A} \in \mathbb{R}^{(U+I) \times (U+I)}$ representing the bidirectional interactions between users and items:

$$\mathbf{A} = \begin{bmatrix} 0 & \mathbf{R} \\ \mathbf{R}^\top & 0 \end{bmatrix}, \quad (1)$$

leading to $\mathcal{G} = \{\mathcal{U} \cup \mathcal{I}, \mathbf{A}\}$ as the user-item bipartite and undirected graph. Moreover, we connote the user- and item-*projected* graphs as $\mathcal{G}_{\mathcal{U}} = \{\mathcal{U}, \mathbf{A}^{\mathcal{U}}\}$ and $\mathcal{G}_{\mathcal{I}} = \{\mathcal{I}, \mathbf{A}^{\mathcal{I}}\}$. In this respect, let $\mathbf{R}^{\mathcal{U}}$ and $\mathbf{R}^{\mathcal{I}}$ be the user-user and item-item interaction matrices:

$$\mathbf{R}^{\mathcal{U}} = \mathbf{R} \cdot \mathbf{R}^\top, \quad \mathbf{R}^{\mathcal{I}} = \mathbf{R}^\top \cdot \mathbf{R}, \quad (2)$$

which indicate the co-occurrences among users and items, respectively. Trivially, the corresponding adjacency matrices $\mathbf{A}^{\mathcal{U}}$ and $\mathbf{A}^{\mathcal{I}}$ are obtained as:

$$\mathbf{A}^{\mathcal{U}} = \mathbf{R}^{\mathcal{U}}, \quad \mathbf{A}^{\mathcal{I}} = \mathbf{R}^{\mathcal{I}}. \quad (3)$$

We use the introduced concepts and notations to describe three topological aspects of the user-item graph and re-interpret them under the lens of recommender systems.

Node degree. Let $\mathcal{N}_u = \{i \mid \mathbf{R}_{u,i} = 1\}$ and $\mathcal{N}_i = \{u \mid \mathbf{R}_{u,i} = 1\}$ be the neighborhood sets for u and i , respectively. By generalizing this definition, let $\mathcal{N}_u^{(l)}$ and $\mathcal{N}_i^{(l)}$ be the sets of neighborhood nodes for user u and item i at l distance hops. Thus, the node degrees for u and i (i.e., $\sigma_u = |\mathcal{N}_u^{(1)}|$ and $\sigma_i = |\mathcal{N}_i^{(1)}|$) represent the number of item and user nodes directly connected with u and i , respectively. The average user and item node degrees are:

$$\sigma_{\mathcal{U}} = \frac{1}{U} \sum_{u \in \mathcal{U}} |\mathcal{N}_u^{(1)}|, \quad \sigma_{\mathcal{I}} = \frac{1}{I} \sum_{i \in \mathcal{I}} |\mathcal{N}_i^{(1)}|. \quad (4)$$

RECSYS RE-INTERPRETATION. *The node degree in the user-item graph stands for the number of items (users) interacted by a user (item). This is related to the cold-start issue in recommendation, where cold users denote low activity on the platform, while cold items are niche products.*

Node degree alone still fails to provide a deeper outlook on the user-item graph. The following topological characteristics, derived from node degree, expand its formulation to other viewpoints.

Clustering coefficient. For each partition in a bipartite graph, it is interesting to recognize clusters of nodes in terms of how their neighborhoods overlap, independently of the respective sizes. Let v

and w be two nodes from the same partition (e.g., user nodes). Their similarity is the intersection over union of their neighborhoods [32]. By evaluating the metric node-wise, we obtain:

$$\gamma_v = \frac{\sum_{w \in \mathcal{N}_v^{(2)}} \gamma_{v,w}}{|\mathcal{N}_v^{(2)}|}, \quad \text{with } \gamma_{v,w} = \frac{|\mathcal{N}_v^{(1)} \cap \mathcal{N}_w^{(1)}|}{|\mathcal{N}_v^{(1)} \cup \mathcal{N}_w^{(1)}|}, \quad (5)$$

where $\mathcal{N}_v^{(2)}$ is the second-order neighborhood set of v . In this case, we leverage the second-order neighborhood because, in a bipartite graph, nodes from the same partition are connected at (multiple of) 2 hops. The average clustering coefficient on \mathcal{U} and \mathcal{I} is:

$$\gamma_{\mathcal{U}} = \frac{1}{U} \sum_{u \in \mathcal{U}} \gamma_u, \quad \gamma_{\mathcal{I}} = \frac{1}{I} \sum_{i \in \mathcal{I}} \gamma_i. \quad (6)$$

RECSYS RE-INTERPRETATION. *High values of the clustering coefficient indicate that there exists a substantial number of co-occurrences among nodes from the same partition. For instance, when considering the user-side formula, the average clustering coefficient increases if several users share most of their interacted items. The intuition aligns with the rationale behind collaborative filtering: two users are likely to show similar preferences when they interact with the same items.*

The clustering coefficient allows the description of broader portions of the user-item graph compared to the semantics conveyed by node degree. Indeed, the measure takes nodes at 2 hops (i.e., user-item-user and item-user-item connections). Nevertheless, we may want to capture properties for even more extended regions of the graph. For this reason, we introduce one last topological characteristic that goes beyond the 2-hop distance among nodes.

Degree assortativity. In real-world graphs, nodes tend to gather when they share similar characteristics. Such a tendency is measured through the assortativity coefficient. Depending on the semantics of “node similarity”, there exist different formulations for assortativity [42]. For the sake of this work, we consider the assortativity coefficient based on the scalar properties of graph nodes, for instance, their degree. Let $\mathcal{D} = \{d_1, d_2, \dots\}$ be the set of unique node degrees in the graph, and let e_{d_h, d_k} be the fraction of edges connecting nodes with degrees d_h and d_k . Then, let q_{d_h} be the probability distribution to choose a node with degree d_h after having selected a node with the same degree (i.e., the *excess* degree distribution). The degree assortativity coefficient is calculated as:

$$\rho = \frac{\sum_{d_h, d_k} d_h d_k (e_{d_h, d_k} - q_{d_h} q_{d_k})}{std_q^2}, \quad (7)$$

where std_q is the standard deviation of the distribution q . Note that, for its formulation, the degree assortativity is similar to a correlation measure (e.g., Pearson correlation). Following the same rationale of the clustering coefficient, we are interested in finding similarity patterns among nodes from the same partition. For this reason, we first apply the projection of the user-item bipartite graph for both users and items to obtain the user- (i.e., $\mathcal{G}_{\mathcal{U}}$) and item- (i.e., $\mathcal{G}_{\mathcal{I}}$) projected graphs. Then, we calculate the degree assortativity coefficients for $\mathcal{G}_{\mathcal{U}}$ and $\mathcal{G}_{\mathcal{I}}$, namely, $\rho_{\mathcal{U}}$ and $\rho_{\mathcal{I}}$.

RECSYS RE-INTERPRETATION. In recommendation, the degree assortativity calculated user- and item-wise represents the tendency of users with the same activity level on the platform and items with the same popularity to gather, respectively. Since we calculate the degree assortativity on the complete user-user and item-item co-occurrence graphs, we deem this characteristic to provide a broader view of the dataset than the clustering coefficient. For this reason, to give an intuition of degree assortativity, we borrow the concept of search space traversal depth in search algorithms theory. That is, we re-interpret degree assortativity in recommendation as a topological characteristic showing a strong look-ahead nature.

3 TOPOLOGICAL CHARACTERISTICS IN GRAPH-BASED RECOMMENDATION

Since GNNs-based recommender systems are designed to work on the bipartite and undirected user-item graph, we seek to understand how and to what extent such models (explicitly) integrate topological properties into their formulations. Thus, we select three popular and recent approaches in GNNs-based recommendation (i.e., LightGCN [25], DGCF [60], and SVD-GCN [44]). Then, we re-formulate their techniques to make the topological data characteristics explicitly emerge.

As additional background with respect to the above, we introduce the notations $\mathbf{e}_u \in \mathbb{R}^b$ and $\mathbf{e}_i \in \mathbb{R}^b$ as the initial embeddings of the nodes for user u and item i , respectively, where $b \ll U, I$. Then, in the case of message-propagation at different layers, we also introduce the notations $\mathbf{e}_u^{(l)}$ and $\mathbf{e}_i^{(l)}$ to indicate the updated node embeddings for user u and item i after l propagation layers, with $0 \leq l \leq L$ (note that $\mathbf{e}_u^{(0)} = \mathbf{e}_u$ and $\mathbf{e}_i^{(0)} = \mathbf{e}_i$).

LightGCN. He et al. [25] propose to lighten the graph convolutional layer presented in Kipf and Welling [30] for the recommendation task. Specifically, their layer removes feature transformation and non-linearities:

$$\mathbf{e}_u^{(l)} = \sum_{i' \in \mathcal{N}_u^{(1)}} \frac{A_{ui'} \mathbf{e}_{i'}^{(l-1)}}{\sqrt{\sigma_u \sigma_{i'}}}, \quad \mathbf{e}_i^{(l)} = \sum_{u' \in \mathcal{N}_i^{(1)}} \frac{A_{iu'} \mathbf{e}_{u'}^{(l-1)}}{\sqrt{\sigma_i \sigma_{u'}}}, \quad (8)$$

where each neighbor contribution is weighted through the corresponding entry in the normalized Laplacian adjacency matrix to flatten the differences among nodes with high and low degrees. Since $A_{ui'} = 1, \forall i' \in \mathcal{N}_u^{(1)}$ (the dual holds for $A_{iu'}$), the contribution weighting comes only from the denominator.

DGCF. Wang et al. [60] assume that user-item interactions are decomposed into a set of independent intents, representing the specific aspects users may be interested in when interacting with items. In this respect, the authors propose to iteratively learn a set of weighted adjacency matrices $\{\tilde{\mathbf{A}}_1, \tilde{\mathbf{A}}_2, \dots\}$, where each of them records the user-item importance weights based on the specific intent it represents. Then, they introduce a graph disentangling layer for each weighted adjacency matrix:

$$\mathbf{e}_{u,*}^{(l)} = \sum_{i' \in \mathcal{N}_u^{(1)}} \frac{\tilde{A}_{ui',*} \mathbf{e}_{i',*}^{(l-1)}}{\sqrt{\sigma_{u,*} \sigma_{i',*}}}, \quad \mathbf{e}_{i,*}^{(l)} = \sum_{u' \in \mathcal{N}_i^{(1)}} \frac{\tilde{A}_{iu',*} \mathbf{e}_{u',*}^{(l-1)}}{\sqrt{\sigma_{i,*} \sigma_{u',*}}}, \quad (9)$$

where $\tilde{A}_{ui',*}$ and $\mathbf{e}_{i',*}^{(l-1)}$ are the learned importance weight of user u on item i' and the embedding of item i' for any intent, while

$\sigma_{u,*}$ is the corresponding node degree calculated on $\tilde{\mathbf{A}}_*$ (the same applies for the item side).

SVD-GCN. Peng et al. [44] propose a reformulation of the GCN-based message-passing which leverages the similarities between graph convolutional layers and singular value decomposition (i.e., SVD). Specifically, they rewrite the message-passing introduced in LightGCN by making two aspects explicitly emerge, namely: (i) even- and odd-connection message aggregations, and (ii) singular values and vectors obtained by decomposing the user-item interaction matrix \mathbf{R} through SVD. On such basis, the authors' assumption is that the traditional graph convolutional layer intrinsically learns a low-rank representation of the user-item interaction matrix where components corresponding to larger singular values tend to be enhanced. They reinterpret the over-smoothing effect as an increasing gap between singular values when stacking more and more layers. The embeddings for users and items are obtained as follows:

$$\mathbf{e}_u = \mathbf{p}_u \exp(a_1 \lambda) \cdot \mathbf{W}, \quad \mathbf{e}_i = \mathbf{q}_i \exp(a_1 \lambda) \cdot \mathbf{W}, \quad (10)$$

where: (i) \mathbf{p}_u and \mathbf{q}_i are the left and right singular vectors of the normalized user-item interaction matrix for user u and item i ; (ii) $\exp(\cdot)$ is the exponential function; (iii) a_1 is a tunable hyper-parameter of the model; (iv) λ is the vector of the largest singular values of the normalized user-item matrix; (v) \mathbf{W} is a trainable matrix to perform feature transformation. Note that the highest singular value λ_{max} and the maximum node degree $\max(\mathcal{D})$ in the user-item interaction matrix are associated by the following inequality:

$$\lambda_{max} \leq \frac{\max(\mathcal{D})}{\max(\mathcal{D}) + a_2}, \quad (11)$$

where a_2 is another tunable hyper-parameter of the model to control the gap among singular values. Moreover, the authors recognize the importance of different types of relationships during the message-passing (i.e., user-item, user-user, item-item). For this reason, they decide to augment the loss function with other components addressing also the similarities among node embeddings from the same partition:

$$\begin{aligned} \min_{\mathbf{e}_v, \mathbf{e}_w, \mathbf{e}_j} & - \sum_{(v,w) \in (\mathbb{R}_s^+)^+} \log(\text{sig}(\mathbf{e}_v^\top \cdot \mathbf{e}_w)) + \\ & - \sum_{(v,j) \in (\mathbb{R}_s^+)^-} \log(\text{sig}(-\mathbf{e}_v^\top \cdot \mathbf{e}_j)), \end{aligned} \quad (12)$$

where v, w , and j are nodes from the same partition, and \mathbb{R}^* is the interaction matrix of that partition.

OBSERVATION. The analyzed GNN models explicitly utilize the node degree information during the representation learning phase, each of them in a different way. However, clustering coefficient and degree assortativity, which share similarities with node degree's semantics, do not seem to have an evident representation within the models' formulations. Under this perspective, our study will also serve to test what topological aspects GNNs-based recommender systems can (un)intentionally capture during their training.

This observation paves the way to a further question: **are (topological) dataset characteristics influencing the recommendation performance of GNNs-based recommender systems?**

4 PROPOSED EVALUATION PIPELINE

To answer such a question, we present our proposed evaluation pipeline to assess the impact of *classical* and *topological* characteristics on the performance of GNNs-based recommender systems. As already done in similar works [2, 18, 37], our goal is to design an explanatory statistical model which finds dependencies between dataset characteristics and recommendation performance.

To this end, the involved pipeline steps are (i) collection of a large pool of recommendation datasets, (ii) calculation of their *classical* and *topological* data characteristics, (iii) train/evaluation of the GNNs-based recommender systems on the recommendation datasets to collect the performance measures (leveraging the same experimental setting as presented in [7]), (iv) design and fit of the explanatory framework on the so-collected data characteristics/recommendation performance samples.

4.1 Step 1: Recommendation data collection

To fit our explanatory framework, it becomes imperative to collect a set of samples (with dataset characteristics and the corresponding models' recommendation performance) that is large enough to ensure the statistical significance of the conducted analysis under a certain confidence threshold. While in principle this process would imply the adoption of several recommendation datasets from the literature, this is in fact infeasible for time and computational reasons because it would require training and evaluating the GNNs-based models on a **very high number of large** datasets.

To this end, and inspired by similar works [2, 18], we propose to select some popular recommendation datasets, and manipulate them through dataset sampling strategies [15, 49, 65] for the generation of several small recommendation datasets to conduct our study effectively. However, given the specific nature of the recommendation models we are dealing with (i.e., GNNs-based recommender systems) and differently from [2, 18], we decide to use ad-hoc **graph** sampling strategies such as node- and edge-dropout, which have gained recent attention in graph learning literature [53, 65]. Indeed, as we vary the dropout rate, **we seek to collect samples that still resemble the original recommendation datasets but can encompass a wide set of topological graph structures.**

As for the original recommendation datasets, we use specific versions of Yelp2018 [44], Gowalla [25], and Amazon-Book [59]. The usage of such datasets is motivated by their popularity in GNNs-based recommendation [22, 25, 33, 41]. Yelp2018 [9] collects data about users and businesses interactions, Amazon-Book is a sub-category of the Amazon dataset [24], and Gowalla [16] is a social-based dataset where users share their locations.

Given the original recommendation dataset, represented as a bipartite and undirected user-item graph, we first randomly select a dropout rate in the range [0.7, 0.9] and a graph sampling strategy in {node-dropout, edge-dropout}. Second, the sampled graph is obtained by either dropping at random nodes or edges depending on the selected dropout rate and sampling strategy. This process is iteratively repeated for each original dataset; for the sake of this work, and following [2, 18], we decide to collect 600 samples from each original data, for a total of 1,800 generated samples.

4.2 Step 2: Characteristics calculation

After the recommendation data collection, the second pipeline step involves the calculation of *classical* and *topological* characteristics of the samples. While we have already focused on the selected *topological* properties, in terms of *classical* characteristics, we select space size, shape, density, and Gini coefficient calculated on the user and item side (once again, we suggest the reader refer to [2, 18] for a formal presentation). Then, inspired by similar works [2, 18], we decide to apply the log10-scale to the formulation of some characteristics to obtain values within comparable order of magnitude, thus making the training of the explanatory model more stable. In the remaining of the paper, for the sake of easy understanding, we use the following notation to indicate the selected data characteristics: $SpaceSize_{log}$, $Shape_{log}$, $Density_{log}$, $Gini-U$, $Gini-I$, $AvgDeg-U_{log}$, $AvgDeg-I_{log}$, $AvgClustC-U_{log}$, $AvgClustC-I_{log}$, $Assort-U$, $Assort-I$, where “log” refers to the log10-scale normalization. Empirically, we note that the vast majority of calculated characteristics appear loosely correlated, hence further supporting their adoption for our analysis (Figure 1).

4.3 Step 3: Train/test GNNs-based models

As a third pipeline step, we train and evaluate the selected GNNs-based recommender systems on the pool of generated datasets. To this end, we follow the same experimental setting proposed in [7], and use Elliot as the reproducibility framework to run our experiments [3, 38]. Thus, we perform the random subsampling strategy to split each dataset into train and test (80% and 20%, respectively). Then, we retain the 10% of the train as validation for the early stopping to avoid overfitting. To train LightGCN, DGCF, and SVD-GCN, we fix their configurations (i.e., hyper-parameters and patience for the early stopping) to the best values according to the original papers, since our scope is not to fine-tune them. Finally, following the literature, we use the Recall@20 calculated on the validation for the early stopping and evaluate the models by assessing the same metric on the test set. To foster the reproducibility of the results, we share the GitHub repository with all codes, datasets, configuration files, and documentation: <https://github.com/sisinflab/Topology-Graph-Collaborative-Filtering>.

4.4 Step 4: Explanatory model

Finally, we aim to fit an explanatory model to the collected dataset characteristics/recommendation performance samples. Statistical models can be utilized to elucidate the relationship between a hypothesized cause of a phenomenon (i.e., independent variables) and its effect (measured through dependent variables). While various potential functions can be used to fit the independent variables to the dependent ones, we opt to utilize a linear regression model for two reasons: (i) to adhere to the same methodology employed in recent studies such as [2, 18], and (ii) to derive explanations on the performance impact of data characteristics through linear dependencies, which represents the most straightforward and intuitive strategy. From this intuition, we use a regression model:

$$y = \epsilon + \theta_0 + \theta_c X_c. \quad (13)$$

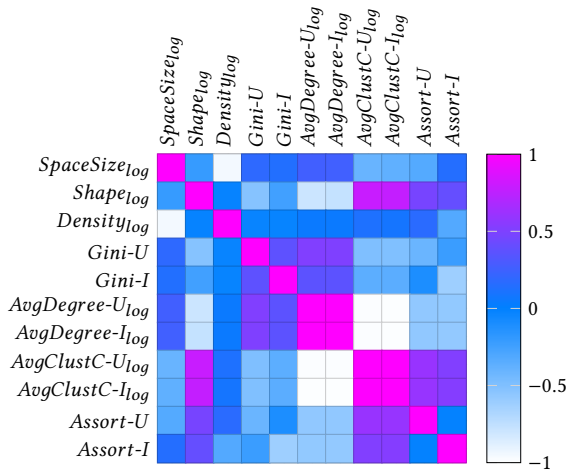


Figure 1: Pearson correlation of the selected characteristics. Many values in $[-0.5, 0.5]$ indicate loosely correlated pairs.

We recall that our goal is to test if the factors related to the data characteristics (i.e., X_c) can explain the effect on the recommendation system’s performance (i.e., y). Therefore, in Equation 13, we denote by $\theta_c = [\theta_1, \dots, \theta_c]$ the vector of regression coefficients each of whom is associated with the c -th feature (data characteristic considered here), $X_c \in \mathbb{R}^{M \times C}$ the matrix containing the data characteristic values for each sample in the training set, and y the vector containing the values of the performance measure associated with all samples in the training set. Moreover, under the assumption of mean-centered data, θ_0 expresses the expected value of y (i.e., in this case, the expected recommendation performance). The regression model is trained through Ordinary Least Squares (OLS):

$$(\theta_0^*, \theta_c^*) = \min_{\theta_0, \theta_c} \frac{1}{2} \|y - \theta_0 - \theta_c X_c\|_2^2. \quad (14)$$

To show how the recommendation performance is related to dataset characteristics, we utilize the basic regression model presented in Equation 14 to maximize the R^2 coefficient. This approach allows us to effectively motivate the impact of the θ_c coefficients on the recommendation system’s effectiveness, as outlined in [20] for any regression model.

5 RESULTS AND DISCUSSION

We aim to answer two research questions: **RQ1**) What is the impact of *classical* and *topological* characteristics on the performance of GNNs-based recommender systems?; **RQ2**) Is the dataset generation through node- and edge-dropout differently influencing the explanations of our model?

5.1 Impact of characteristics (RQ1)

We assess the impact of *classical* and *topological* characteristics on the accuracy performance (i.e., Recall@20) of GNNs-based recommender systems. Figure 2 displays, for each dataset/GNNs-based model setting, the learned regression coefficients of the explanatory framework. Specifically, each bar plot indicates the relative impact of a given data characteristic on the Recall@20 performance;

moreover, bar plot length and direction represent the magnitude of this impact and whether there exists a direct/inverse correspondence between characteristic and performance. Finally, to assess the goodness of the results, we also estimate the adj. (*adjusted*) R^2 [2, 18] of the regression model, along with the statistical significance of the learned coefficients (the darker the bar plots, the higher the statistical significance). Overall, the adj. R^2 is above 95%, proving the ability of the regression model to explain the accuracy recommendation performance through the measured characteristics. Hereinafter, we further decompose the regression results by categorizing the characteristics into *classical* and *topological*.

5.1.1 Classical characteristics. Previous works [2] have assessed the impact of *classical* characteristics on neighbor- and factorization-based models separately. However, a careful search of the relevant literature yields that no study has investigated whether such characteristics influence GNNs-based recommender systems [37]; as already discussed, GNNs-based models involve both a neighborhood- and factorization-based module, where the former corresponds to the message-passing, and the latter to the users’ and items’ latent factors as in traditional latent-based recommendation [46, 47].

Interestingly, Figure 2 suggests that, in GNNs-based recommender systems, the factorization component might be more influential than the neighborhood one on the recommendation performance. In this respect, if we refer to the results from [2], we observe that factorization- and GNNs-based approaches are particularly aligned, considering: (i) the inverse correspondence between the performance metric and the $Shape_{log}$ in almost all settings, meaning that when the number of users is higher than the number of items in the system, performance may decrease; (ii) the direct correspondence between the performance metric and $Density_{log}$ and $Gini-I$. As for (ii), the density is historically known as one of the core problems in recommendation (i.e., data sparsity), so it becomes evident why also GNNs-based recommender systems’ performance benefits from denser (i.e., less sparse) datasets. The Gini index measures the dissimilar distribution of items’ interactions in the system and could be related to the tendency of recommender systems to promote popular items from the catalog. That is, when there exist items that have been experienced more frequently than others, both GNNs- and factorization-based recommender systems may be biased towards popular items, and so their accuracy performance increases. Noteworthy, all such observations are supported by the statistical significance of the results.

5.1.2 Topological characteristics. GNNs-based models interpret the user-item data as a graph. Consequently, our novel evaluation analyzes the influence of *topological* characteristics of the selected recommendation datasets on the recommendation performance.

The most evident outcome is that $AvgDegree-U_{log}$ and $AvgDegree-I_{log}$ show a direct correspondence with performance in almost all settings. Indeed, this analytically confirms the intuitions provided in the literature on the recommendation scenario [7] or for other tasks [61], as well as what we already observed in Section 3 regarding the explicit presence of the node degree in the formulations of all the selected GNNs-based approaches. In practical terms, when GNNs-based models are trained on datasets with several interactions for users and items, they learn accurate users’ preferences since each node receives the contribution of numerous neighbor




Figure 2: Visual representation of the impact of dataset characteristics on the recommendation performance (Recall@20) of GNNs-based recommender systems, for each dataset/model setting. Bar plot length and direction represent the impact magnitude and whether there is a direct/inverse correspondence between characteristic and performance. Finally, the darker the bar plots, the higher their statistical significance.

nodes. It is worth noticing that, in absolute values, $AvgDegree-U_{log}$ is more influential than $AvgDegree-I_{log}$ on the overall performance. Hence, under the same average degree gain, a user average degree improvement is preferable since it would lead to better performance.

As far as clustering coefficient and degree assortativity are concerned, we assess how similarities among nodes from the same partition in the graph may impact the recommendation accuracy performance of models. In terms of $AvgClustC-U_{log}$ and $AvgClustC-I_{log}$, the results prove again a strong direct correspondence in almost all settings of GNNs-based models and datasets. Differently from the average degree scenario, the relative importance of the user-side values is much higher than the one of the item-side for LightGCN and DGCF, while the gap sometimes gets narrower in the case of SVD-GCN. This may happen because while LightGCN and DGCF only leverage user-item types of interactions, SVD-GCN also embed the information conveyed in the user- and item-projected graphs in its formulation, thus flattening the different influence of the user-side characteristics over the item-side counterpart.

Interestingly, the $Assort-U$ and $Assort-I$ characteristics exhibit a direct and inverse correspondence to the recommendation metric, respectively. Furthermore, models such as LightGCN and DGCF have slightly larger coefficients for both $Assort-U$ and $Assort-I$ than SVD-GCN. Again, these results have a mathematical justification. Indeed, the strong *lookahead* nature of the assortativity measures (refer again to Section 2) seems to be captured by the multi-layer message-passing performed by LightGCN and DGCF. Conversely, in the case of SVD-GCN, they are less influential, probably because the model acts on the singular values of the adjacency matrix with the effect of limiting the graph convolutional layers' depth to avoid over-smoothing. However, the assortativity results are less significant than the others, so we plan to further investigate this aspect.

SUMMARY. *The results show that: (i) interestingly, the factorization component of GNNs-based recommender systems seems to be impacting the recommendation performance more than the neighborhood (message-passing) component; (ii) while confirming its influence on the recommendation performance, node degree seems not to be a key topological characteristic to distinguish among the different GNNs-based models; indeed, the wider perspective provided by clustering coefficient and (especially) degree assortativity may help to recognize how the different models address the topological properties of the graph, even with unexpected outcomes.*

5.2 Influence of node- and edge-dropout (RQ2)

The current section investigates the influence of node- and edge-dropout on the explanatory model. In the interest of space, we report an extensive analysis of the largest dataset, Gowalla, by considering the performance of SVD-GCN, the most recent GNNs-based approach among the selected ones. To answer the RQ, we provide both a *theoretical* and *analytical* intuition.

5.2.1 Theoretical intuition. Figure 3 displays the relation between the probability distribution of node degrees in the original graph and their degree values on the Gowalla dataset. As evident, high-degree nodes are less popular than low-degree ones, and this resembles the tendency of real-world networks to be *scale-free* [45]. To




Figure 3: Node degree probability distribution on Gowalla. The black points (i.e., the real data) would be approximated by a function in-between the power-law and the exponential.

be more precise, the actual degree probability distribution approximates neither the *power-law* (i.e., representing *scale-free* networks, in green), nor the *exponential* function (i.e., in red), but it would be approximated by a function in-between. Thus, high-degree nodes are even less frequent than they usually are in *scale-free* networks.

The figure helps us understand what the different impacts of node- and edge-dropout on such a graph topology might be. By recalling that node-dropout works by removing nodes (and all the edges connected to them) and edge-dropout eliminates edges and the consequently disconnected nodes, in the worst-case scenario, node-dropout would remove many high-degree nodes from the graph; instead, edge-dropout would eliminate all the edges connected to several nodes and thus disconnect them from the graph. Hence, on average, node-dropout would drop larger portions of the user-item graph than edge-dropout, possibly undermining the goodness of the explanations produced by our explanatory framework. This intuition *theoretically* justifies the *joint* adoption of node- and edge-dropout for the experiments presented in RQ1.

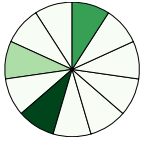
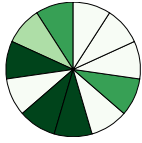
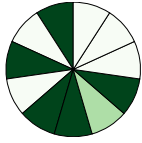
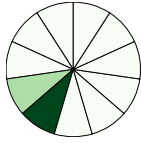
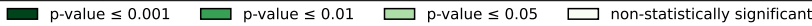
5.2.2 Analytical intuition. To further analytically test the previous intuition, we build four versions of the dataset X_c for the explanatory framework, each obtained with varying portions of node- and edge-dropout, respectively (refer again to Equation (13)). Specifically, the number of samples in X_c changes in accordance to:

$$|X_c| = (1 - \alpha)|X_c^n| + \alpha|X_c^e|, \quad (15)$$

where X_c^n and X_c^e indicate the portion of X_c sampled through node- and edge-dropout, while α is a parameter to control the number of samples from X_c^n and X_c^e contributing to the final dataset X_c . We use $|\cdot|$ as a necessary notation abuse to refer to any dataset size in a simple way. We let α range in $\{0.0, 0.3, 0.7, 1.0\}$, where extreme values of α are used to build the dataset through either node- or edge-dropout; the others combine the two sampling strategies.

Table 1 provides a visual representation of how results vary with different α values. Note that we only report the average sampling statistics and the measures indicating the goodness of the generated explanations (i.e., the calculated $\text{adj}R^2$, and statistical significance of the learned regression coefficients represented in

Table 1: Explanatory framework results with varying values of α on Gowalla for SVD-GCN. In the first row, we provide a visual intuition for the α value changes in all four settings. In the second row, we report the average sampling statistics of the datasets. Finally, in the third row, we indicate the $\text{adj.}R^2$ regression measure, and a pie chart indicating, for each learned regression coefficient, its statistical significance (the darker the more statistically significant).

Node ●●● Edge ○○○	Node ●●○ Edge ●○○	Node ●○○ Edge ●●○	Node ○○○ Edge ●●●
<i>Avg. Stats.</i> Users: 5,828 Items: 7,887 Interactions: 45,620	<i>Avg. Stats.</i> Users: 12,744 Items: 17,229 Interactions: 97,785	<i>Avg. Stats.</i> Users: 21,730 Items: 29,316 Interactions: 160,919	<i>Avg. Stats.</i> Users: 28,526 Items: 38,467 Interactions: 209,659
$\text{adj.}R^2 = 0.745$ 	$\text{adj.}R^2 = 0.969$ 	$\text{adj.}R^2 = 0.987$ 	$\text{adj.}R^2 = 0.991$ 
			

shades of green once again) as this is the only information we need to answer the RQ. In alignment with the *theoretical* intuition, the average sampling statistics show that node-dropout generally retains smaller portions of the graph than the edge-dropout. Then, the regression results highlight that the optimal trade-off between high $\text{adj.}R^2$ and statistical significance of the learned coefficients is reached when combining samples generated through both node- and edge-dropout. On the contrary, the settings with either node- or edge-dropout do not offer the conditions for the regression model to learn meaningful dependencies. This *analytically* justifies the *joint* adoption of node- and edge-dropout for RQ1.

SUMMARY. *The theoretical and analytical evaluation of the explanatory model for different settings of node- and edge-dropout indicates that their joint combination (i.e., the strategy we followed in RQ1) is beneficial to produce meaningful explanations.*

6 CONCLUSION AND FUTURE WORK

In this work, we propose a novel evaluation perspective on graph neural networks (GNNs)-based recommender systems, that aims to assess the influence of the user-item graph topological properties on the recommendation performance of such models. First, we select classical and topological properties of the recommendation data, as well as three popular and recent GNNs-based recommender systems. On such a basis, we design a novel evaluation pipeline, involving: (i) the generation of a large pool of reduced-size recommendation data (sampled from the selected datasets through node- and edge-dropout) that encompass a wide range of topological structures; (ii) the calculation of their dataset characteristics and (iii) evaluation of the models' recommendation performance; (iv) the design and fit of an explanatory model that finds linear dependencies between dataset characteristics and performance. Results, validated by statistical tests and under different sampling

settings, largely demonstrate the presence of strong characteristics/performance correspondences, offering a novel perspective on GNNs-based recommendation. We plan to extend our investigation to other GNNs-based recommender systems and consider additional recommendation metrics (e.g., bias and fairness dimensions). Moreover, driven by the derived insights, we aim to propose a novel GNNs-based approach that exploits the topological structure of the graph to eventually boost the recommendation performance.

ACKNOWLEDGMENTS

The authors acknowledge the partial support by: CT_FINCONS_III, OVS Fashion Retail Reloaded, LUTECH DIGITALE 4.0, Secure Safe Apulia, IDENTITA, REACH-XY. The authors acknowledge the CINECA award under the ISCRA initiative for the availability of high-performance computing resources. This work has been carried out while A. C. M. Mancino was enrolled in the Italian National Doctorate on Artificial Intelligence run by the Sapienza University of Rome in collaboration with the Polytechnic University of Bari.

REFERENCES

- [1] Sami Abu-El-Hajja, Bryan Perozzi, Amol Kapoor, Nazanin Alipourfard, Kristina Lerman, Hrayr Harutyunyan, Greg Ver Steeg, and Aram Galstyan. 2019. MixHop: Higher-Order Graph Convolutional Architectures via Sparsified Neighborhood Mixing. In *ICML (Proceedings of Machine Learning Research, Vol. 97)*. PMLR, 21–29.
- [2] Gediminas Adomavicius and Jingjing Zhang. 2012. Impact of data characteristics on recommender systems performance. *ACM Trans. Manag. Inf. Syst.* 3, 1 (2012), 3:1–3:17.
- [3] Vito Walter Anelli, Alejandro Bellogín, Antonio Ferrara, Daniele Malitesta, Felice Antonio Merra, Claudio Pomo, Francesco Maria Donini, and Tommaso Di Noia. 2021. Elliot: A Comprehensive and Rigorous Framework for Reproducible Recommender Systems Evaluation. In *SIGIR*. ACM, 2405–2414.
- [4] Vito Walter Anelli, Yashar Deldjoo, Tommaso Di Noia, Daniele Malitesta, Vincenzo Paparella, and Claudio Pomo. 2023. Auditing Consumer- and Producer-Fairness in Graph Collaborative Filtering. In *ECIR (1) (Lecture Notes in Computer Science, Vol. 13980)*. Springer, 33–48.
- [5] Vito Walter Anelli, Yashar Deldjoo, Tommaso Di Noia, Eugenio Di Sciascio, Antonio Ferrara, Daniele Malitesta, and Claudio Pomo. 2022. How Neighborhood Exploration influences Novelty and Diversity in Graph Collaborative Filtering. In *MORS@RecSys (CEUR Workshop Proceedings, Vol. 3268)*. CEUR-WS.org.

- [6] Vito Walter Anelli, Yashar Deldjoo, Tommaso Di Noia, Eugenio Di Sciascio, Antonio Ferrara, Daniele Malitesta, and Claudio Pomo. 2022. Reshaping Graph Recommendation with Edge Graph Collaborative Filtering and Customer Reviews. In *DL4SR@CIKM (CEUR Workshop Proceedings, Vol. 3317)*. CEUR-WS.org.
- [7] Vito Walter Anelli, Daniele Malitesta, Claudio Pomo, Alejandro Bellogin, Eugenio Di Sciascio, and Tommaso Di Noia. 2023. Challenging the Myth of Graph Collaborative Filtering: a Reasoned and Reproducibility-driven Analysis. In *RecSys*. ACM, 350–361.
- [8] Vito Walter Anelli, Tommaso Di Noia, Eugenio Di Sciascio, Antonio Ferrara, and Alberto Carlo Maria Mancino. 2021. Sparse Feature Factorization for Recommender Systems with Knowledge Graphs. In *RecSys*. ACM, 154–165.
- [9] Nabihah Asghar. 2016. Yelp Dataset Challenge: Review Rating Prediction. *CoRR* abs/1605.05362 (2016).
- [10] Bahare Askari, Jaroslaw Szlichta, and Amirali Salehi-Abari. 2021. Variational Autoencoders for Top-K Recommendation with Implicit Feedback. In *SIGIR*. ACM, 2061–2065.
- [11] Xuheng Cai, Chao Huang, Lianghao Xia, and Xubin Ren. 2023. LightGCL: Simple Yet Effective Graph Contrastive Learning for Recommendation. In *ICLR*. OpenReview.net.
- [12] Daniele Castellana and Federico Errica. 2023. Investigating the Interplay between Features and Structures in Graph Learning. *CoRR* abs/2308.09570 (2023).
- [13] Deli Chen, Yankai Lin, Wei Li, Peng Li, Jie Zhou, and Xu Sun. 2020. Measuring and Relieving the Over-Smoothing Problem for Graph Neural Networks from the Topological View. In *AAAI*. AAAI Press, 3438–3445.
- [14] Lei Chen, Le Wu, Richang Hong, Kun Zhang, and Meng Wang. 2020. Revisiting Graph Based Collaborative Filtering: A Linear Residual Graph Convolutional Network Approach. In *AAAI*. AAAI Press, 27–34.
- [15] Ting Chen, Yizhou Sun, Yue Shi, and Liangjie Hong. 2017. On Sampling Strategies for Neural Network-based Collaborative Filtering. In *KDD*. ACM, 767–776.
- [16] Eunjoon Cho, Seth A. Myers, and Jure Leskovec. 2011. Friendship and mobility: user movement in location-based social networks. In *KDD*. ACM, 1082–1090.
- [17] Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep Neural Networks for YouTube Recommendations. In *RecSys*. ACM, 191–198.
- [18] Yashar Deldjoo, Tommaso Di Noia, Eugenio Di Sciascio, and Felice Antonio Merri. 2020. How Dataset Characteristics Affect the Robustness of Collaborative Recommendation Models. In *SIGIR*. ACM, 951–960.
- [19] Antonio Ferrara, Vito Walter Anelli, Alberto Carlo Maria Mancino, Tommaso Di Noia, and Eugenio Di Sciascio. 2023. KGFlex: Efficient Recommendation with Sparse Feature Factorization and Knowledge Graphs. *Trans. Recomm. Syst.* 1, 4 (2023), 1–30.
- [20] James Gareth Witten Daniela, Hastie Trevor, and Tibshirani Robert. 2013. *An introduction to statistical learning: with applications in R*. Springer.
- [21] Jhony H. Giraldo, Konstantinos Skianis, Thierry Bouwmans, and Fragkiskos D. Malliaros. 2023. On the Trade-off between Over-smoothing and Over-squashing in Deep Graph Neural Networks. In *CIKM*. ACM, 566–576.
- [22] Kaiqi Gong, Xiao Song, Senzhang Wang, Songsong Liu, and Yong Li. 2022. ITSM-GCN: Informative Training Sample Mining for Graph Convolutional Network-based Collaborative Filtering. In *CIKM*. ACM, 614–623.
- [23] William L. Hamilton, Zhitaoying, and Jure Leskovec. 2017. Inductive Representation Learning on Large Graphs. In *NIPS*. 1024–1034.
- [24] Ruining He and Julian J. McAuley. 2016. Ups and Downs: Modeling the Visual Evolution of Fashion Trends with One-Class Collaborative Filtering. In *WWW*. ACM, 507–517.
- [25] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yong-Dong Zhang, and Meng Wang. 2020. LightGCN: Simplifying and Powering Graph Convolution Network for Recommendation. In *SIGIR*. ACM, 639–648.
- [26] Xiangnan He, Zhankui He, Xiaoyu Du, and Tat-Seng Chua. 2018. Adversarial Personalized Ranking for Recommendation. In *SIGIR*. ACM, 355–364.
- [27] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural Collaborative Filtering. In *WWW*. ACM, 173–182.
- [28] Chao Huang, Lianghao Xia, Xiang Wang, Xiangnan He, and Dawei Yin. 2022. Self-Supervised Learning for Recommendation. In *CIKM*. ACM, 5136–5139.
- [29] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschot, Ce Liu, and Dilip Krishnan. 2020. Supervised Contrastive Learning. In *NeurIPS*.
- [30] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR (Poster)*. OpenReview.net.
- [31] Johannes Klicpera, Aleksandar Bojchevski, and Stephan Gunnemann. 2019. Predict then Propagate: Graph Neural Networks meet Personalized PageRank. In *ICLR (Poster)*. OpenReview.net.
- [32] Matthieu Latapy, Clémence Magnien, and Nathalie Del Vecchio. 2008. Basic notions for the analysis of large two-mode networks. *Soc. Networks* 30, 1 (2008), 31–48.
- [33] Zihan Lin, Changxin Tian, Yupeng Hou, and Wayne Xin Zhao. 2022. Improving Graph Collaborative Filtering with Neighborhood-enriched Contrastive Learning. In *WWW*. ACM, 2320–2329.
- [34] Siwei Liu, Iadh Ounis, and Craig Macdonald. 2022. An MLP-based Algorithm for Efficient Contrastive Graph Recommendations. In *SIGIR*. ACM, 2431–2436.
- [35] Sitao Luan, Mingde Zhao, Xiao-Wen Chang, and Doina Precup. 2019. Break the Ceiling: Stronger Multi-scale Deep Graph Convolutional Networks. In *NeurIPS*. 10943–10953.
- [36] Jianxin Ma, Peng Cui, Kun Kuang, Xin Wang, and Wenwu Zhu. 2019. Disentangled Graph Convolutional Networks. In *ICML (Proceedings of Machine Learning Research, Vol. 97)*. PMLR, 4212–4221.
- [37] Daniele Malitesta, Claudio Pomo, Vito Walter Anelli, Alberto Carlo Maria Mancino, Eugenio Di Sciascio, and Tommaso Di Noia. 2023. A Topology-aware Analysis of Graph Collaborative Filtering. *CoRR* abs/2308.10778 (2023).
- [38] Daniele Malitesta, Claudio Pomo, Vito Walter Anelli, Tommaso Di Noia, and Antonio Ferrara. 2023. An Out-of-the-Box Application for Reproducible Graph Collaborative Filtering extending the Elliot Framework. In *UMAP (Adjunct Publication)*. ACM, 12–15.
- [39] Daniele Malitesta, Claudio Pomo, and Tommaso Di Noia. 2023. Graph Neural Networks for Recommendation: Reproducibility, Graph Topology, and Node Representation. *CoRR* abs/2310.11270 (2023).
- [40] Alberto Carlo Maria Mancino, Antonio Ferrara, Salvatore Bui, Daniele Malitesta, Tommaso Di Noia, and Eugenio Di Sciascio. 2023. KGTORE: Tailored Recommendations through Knowledge-aware GNN Models. In *RecSys*. ACM, 576–587.
- [41] Kelong Mao, Jieming Zhu, Xi Xiao, Biao Lu, Zhaowei Wang, and Xiuqiang He. 2021. UltraGCN: Ultra Simplification of Graph Convolutional Networks for Recommendation. In *CIKM*. ACM, 1253–1262.
- [42] M. E. J. Newman. 2003. Mixing patterns in networks. *Phys. Rev. E* 67 (Feb 2003), 026126. Issue 2. <https://doi.org/10.1103/PhysRevE.67.026126>
- [43] Bibek Paudel, Fabian Christoffel, Chris Newell, and Abraham Bernstein. 2017. Updatable, Accurate, Diverse, and Scalable Recommendations for Interactive Applications. *ACM Trans. Interact. Syst.* 7, 1 (2017), 1:1–1:34.
- [44] Shaowen Peng, Kazunari Sugiyama, and Tsunenori Mine. 2022. SVD-GCN: A Simplified Graph Convolution Paradigm for Recommendation. In *CIKM*. ACM, 1625–1634.
- [45] Erzsébet Ravasz and Albert-László Barabási. 2003. Hierarchical organization in complex networks. *Phys. Rev. E* 67 (Feb 2003), 026112. Issue 2. <https://doi.org/10.1103/PhysRevE.67.026112>
- [46] Steffen Rendle. 2010. Factorization Machines. In *ICDM*. IEEE Computer Society, 995–1000.
- [47] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *UAI*. AUAI Press, 452–461.
- [48] Yu Rong, Wenbing Huang, Tingyang Xu, and Junzhou Huang. 2020. DropEdge: Towards Deep Graph Convolutional Networks on Node Classification. In *ICLR*. OpenReview.net.
- [49] Naveen Sachdeva, Carole-Jean Wu, and Julian J. McAuley. 2022. On Sampling Collaborative Filtering Datasets. In *WSDM*. ACM, 842–850.
- [50] Badrul Munir Sarwar, George Karypis, Joseph A. Konstan, and John Riedl. 2001. Item-based collaborative filtering recommendation algorithms. In *WWW*. ACM, 285–295.
- [51] Yifei Shen, Yongji Wu, Yao Zhang, Caihua Shan, Jun Zhang, Khaled B. Letaief, and Dongsheng Li. 2021. How Powerful is Graph Convolution for Recommendation?. In *CIKM*. ACM, 1619–1629.
- [52] Ilya Shenbin, Anton Alekseev, Elena Tutubalina, Valentin Malykh, and Sergey I. Nikolenko. 2020. RecVAE: A New Variational Autoencoder for Top-N Recommendations with Implicit Feedback. In *WSDM*. ACM, 528–536.
- [53] Juan Shu, Bowei Xi, Yu Li, Fan Wu, Charles A. Kamhoua, and Jianzhu Ma. 2022. Understanding Dropout for Graph Neural Networks. In *WWW (Companion Volume)*. ACM, 1128–1138.
- [54] Jake Topping, Francesco Di Giovanni, Benjamin Paul Chamberlain, Xiaowen Dong, and Michael M. Bronstein. 2022. Understanding over-squashing and bottlenecks on graphs via curvature. In *ICLR*. OpenReview.net.
- [55] Rianne van den Berg, Thomas N. Kipf, and Max Welling. 2017. Graph Convolutional Matrix Completion. *CoRR* abs/1706.02263 (2017).
- [56] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. In *ICLR (Poster)*. OpenReview.net.
- [57] Jun Wang, Lantao Yu, Weinan Zhang, Yu Gong, Yinghui Xu, Benyou Wang, Peng Zhang, and Dell Zhang. 2017. IRGAN: A Minimax Game for Unifying Generative and Discriminative Information Retrieval Models. In *SIGIR*. ACM, 515–524.
- [58] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. 2019. Neural Graph Collaborative Filtering. In *SIGIR*. ACM, 165–174.
- [59] Xiang Wang, Tinglin Huang, Dingxian Wang, Yancheng Yuan, Zhenguang Liu, Xiangnan He, and Tat-Seng Chua. 2021. Learning Intents behind Interactions with Knowledge Graph for Recommendation. In *WWW*. ACM / IW3C2, 878–887.
- [60] Xiang Wang, Hongye Jin, An Zhang, Xiangnan He, Tong Xu, and Tat-Seng Chua. 2020. Disentangled Graph Collaborative Filtering. In *SIGIR*. ACM, 1001–1010.
- [61] Yu Wang, Tong Zhao, Yuying Zhao, Yunchao Liu, Xueqi Cheng, Neil Shah, and Tyler Derr. 2023. A Topological Perspective on Demystifying GNN-Based Link Prediction Performance. *CoRR* abs/2310.04612 (2023).

- [62] Zhenyi Wang, Huan Zhao, and Chuan Shi. 2022. Profiling the Design Space for Graph Neural Networks based Collaborative Filtering. In *WSDM*. ACM, 1109–1119.
- [63] Lanning Wei, Huan Zhao, and Zhiqiang He. 2022. Designing the Topology of Graph Neural Networks: A Novel Feature Fusion Perspective. In *WWW*. ACM, 1381–1391.
- [64] Felix Wu, Amauri H. Souza Jr., Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Q. Weinberger. 2019. Simplifying Graph Convolutional Networks. In *ICML (Proceedings of Machine Learning Research, Vol. 97)*. PMLR, 6861–6871.
- [65] Jiancan Wu, Xiang Wang, Fuli Feng, Xiangnan He, Liang Chen, Jianxun Lian, and Xing Xie. 2021. Self-supervised Graph Learning for Recommendation. In *SIGIR*. ACM, 726–735.
- [66] Shiwen Wu, Fei Sun, Wentao Zhang, Xu Xie, and Bin Cui. 2023. Graph Neural Networks in Recommender Systems: A Survey. *ACM Comput. Surv.* 55, 5 (2023), 97:1–97:37.
- [67] Wei Wu, Chao Wang, Dazhong Shen, Chuan Qin, Liyi Chen, and Hui Xiong. 2024. AFDGCF: Adaptive Feature De-correlation Graph Collaborative Filtering for Recommendations. *CoRR* abs/2403.17416 (2024).
- [68] Keyulu Xu, Chengtao Li, Yonglong Tian, Tomohiro Sonobe, Ken-ichi Kawarabayashi, and Stefanie Jegelka. 2018. Representation Learning on Graphs with Jumping Knowledge Networks. In *ICML (Proceedings of Machine Learning Research, Vol. 80)*. PMLR, 5449–5458.
- [69] Liang Yang, Wenmiao Zhou, Weihang Peng, Bingxin Niu, Junhua Gu, Chuan Wang, Xiaochun Cao, and Dongxiao He. 2022. Graph Neural Networks Beyond Compromise Between Attribute and Topology. In *WWW*. ACM, 1127–1135.
- [70] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L. Hamilton, and Jure Leskovec. 2018. Graph Convolutional Neural Networks for Web-Scale Recommender Systems. In *KDD*. ACM, 974–983.
- [71] Junliang Yu, Xin Xia, Tong Chen, Lizhen Cui, Nguyen Quoc Viet Hung, and Hongzhi Yin. 2024. XSimGCL: Towards Extremely Simple Graph Contrastive Learning for Recommendation. *IEEE Trans. Knowl. Data Eng.* 36, 2 (2024), 913–926.
- [72] Junliang Yu, Hongzhi Yin, Xin Xia, Tong Chen, Lizhen Cui, and Quoc Viet Hung Nguyen. 2022. Are Graph Augmentations Necessary?: Simple Graph Contrastive Learning for Recommendation. In *SIGIR*. ACM, 1294–1303.
- [73] Yiding Zhang, Chaozhuo Li, Xing Xie, Xiao Wang, Chuan Shi, Yuming Liu, Hao Sun, Liangjie Zhang, Weiwei Deng, and Qi Zhang. 2022. Geometric Disentangled Collaborative Filtering. In *SIGIR*. ACM, 80–90.
- [74] Lingxiao Zhao and Leman Akoglu. 2020. PairNorm: Tackling Oversmoothing in GNNs. In *ICLR*. OpenReview.net.
- [75] Jiong Zhu, Yujun Yan, Lingxiao Zhao, Mark Heimann, Leman Akoglu, and Danai Koutra. 2020. Beyond Homophily in Graph Neural Networks: Current Limitations and Effective Designs. In *NeurIPS*.